

Performance Specifications


Arvind
Computer Science & Artificial Intelligence Lab
Massachusetts Institute of Technology

March 5, 2008

<http://csg.csail.mit.edu/6.375>

L11-1

Simple processor pipeline



- ◆ Functional behavior is well understood
- ◆ Intuition about performance is lacking
 - Should the branch be resolved in the Decode or Execute stage?
 - Should the branch target address be latched before its use?
- ◆ Experimentation is required to evaluate design alternatives

cycle time?
area?
execution time?

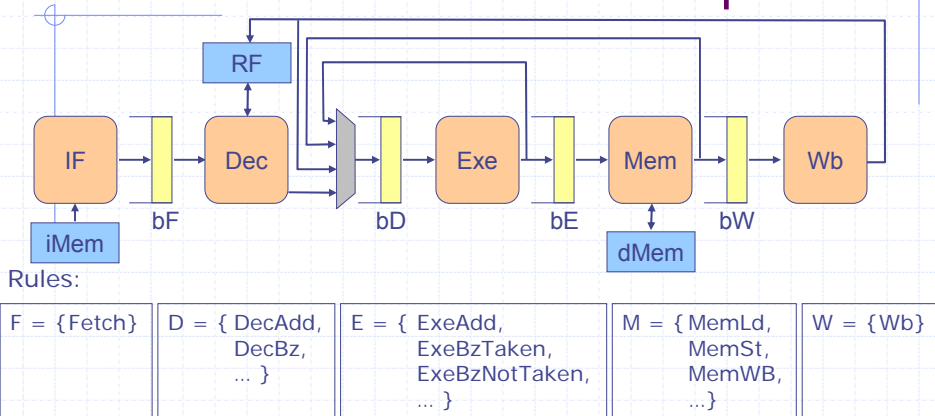
We present a design flow that makes such experimentation easy for the designer

March 5, 2008

<http://csg.csail.mit.edu/6.375>

L11-2

Need for Performance Specs



- What is the design's performance / throughput?
- Reference model implies one rule per cycle execution

Designer's goal is usually different and based on the application!

March 5, 2008

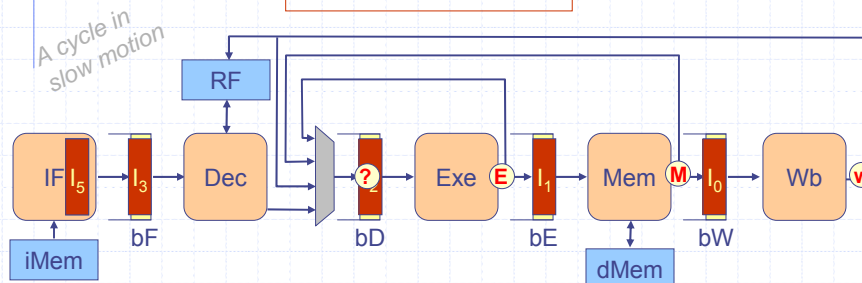
<http://csg.csail.mit.edu/6.375>

L11-3

Pipelining via Performance specification

- ◆ The designer wants a pipeline which executes one instruction every cycle
- ◆ Performance spec for a pipelined processor:

$$W < M < E < D < F$$



March 5, 2008

<http://csg.csail.mit.edu/6.375>

L11-4

More Performance Specification

F = {Fetch}	D = { DecAdd, DecBz, ... }	E = { ExeAdd, ExeBzTaken, ExeBzNotTaken, ... }	M = { MemLd, MemSt, MemWB, ... }	W = {Wb}
-------------	----------------------------------	---	---	----------

We allow the designer to specify performance!

$W < M < E < D < F \quad \equiv \textit{pipelined}$

- 1) $W < M < E^* < D < F \quad \equiv \textit{pipelined except for ExeBzTaken}$
- 2) $W < M < \textit{ExeBzTaken}$

What do the following mean?

$F < D < E < M < W \quad \equiv \textit{unpipelined (assuming buffers start empty)}$

$W < W < M < M < E < E < D < D < F < F \quad \equiv \textit{two-way superscalar!}$

Synthesis algorithms ensure that performance specs are satisfied and guarantee that functionality is not altered.

March 5, 2008

<http://csg.csail.mit.edu/6.375>

L11-5

Why is functionality maintained?

- ◆ A few observations about rule-based systems:
 - Adding a new rule to a system can only introduce new behaviors
 - If the new rule is a *derived* rule, then it does not add new behaviors

◆ Composed rules:

- Given rules:

$R_a: \text{when } \pi_a(s) \Rightarrow s := \delta_a(s);$

$R_b: \text{when } \pi_b(s) \Rightarrow s := \delta_b(s);$

- The composed rule is a derived rule:


$R_{a,b}: \text{when } \pi_a(s) \ \& \ \pi_b(\delta_a(s)) \Rightarrow s := \delta_b(\delta_a(s));$

March 5, 2008

<http://csg.csail.mit.edu/6.375>

L11-6

Scheduling Specifications



```
rule fetch_and_decode (!stallfunc(instr, bu));
    bu.enq(newIt(instr, rf));
    pc <= predIa;
endrule
```

```
rule execAdd
    (it matches tagged EAdd{dst:.rd,src1:.va,src2:.vb});
    rf.upd(rd, va+vb); bu.deq(); endrule
rule execBzTaken(it matches tagged Bz {cond:.cv,addr:.av}
    &&& (cv == 0));
    pc <= av; bu.clear(); endrule
rule execBzNotTaken(it matches tagged Bz {cond:.cv,addr:.av}
    &&& !(cv == 0));
    bu.deq(); endrule
rule execLoad(it matches tagged ELoad{dst:.rd,addr:.av});
    rf.upd(rd, dMem.read(av)); bu.deq(); endrule
rule execStore(it matches tagged EStore{val:.v,addr:.av});
    dMem.write(av, vv); bu.deq(); endrule
```

execAdd < fetch

execBzTaken < fetch

execBzNotTaken < fetch ?

execLoad < fetch


execStore < fetch

March 5, 2008

<http://csg.csail.mit.edu/6.375>

L11-7

Implications for modules



```
rule fetch_and_decode (!stallfunc(instr, bu));
    bu.enq(newIt(instr, rf));
    pc <= predIa;
endrule
```

```
rule execAdd
    (it matches tagged EAdd{dst:.rd,src1:.va,src2:.vb});
    rf.upd(rd, va+vb); bu.deq();
endrule
```

◆ execAdd < fetch ⇒


- rf: sub > upd
- bu: {find, enq} > {first, deq}

March 5, 2008

<http://csg.csail.mit.edu/6.375>

L11-8

Branch rules



```
rule fetch_and_decode (!stallfunc(instr, bu));
    bu.enq(newIt(instr, rf));
    pc <= predIa;
endrule
```

```
rule execBzTaken(it matches tagged Bz {cond:.cv,addr:.av}
    &&& (cv == 0));
    pc <= av; bu.clear(); endrule
```

```
rule execBzNotTaken(it matches tagged Bz {cond:.cv,addr:.av}
    &&& !(cv == 0));
    bu.deq(); endrule
```


- ◆ execBzTaken < fetch ?
 - Should be treated as conflict – give priority to execBzTaken
- ◆ execBzNotTaken < fetch
 - bu: {first, deq} < {find, enq}

March 5, 2008

<http://csg.csail.mit.edu/6.375>

L11-9

Load-Store Rules



```
rule fetch_and_decode (!stallfunc(instr, bu));
    bu.enq(newIt(instr, rf));
    pc <= predIa;
endrule
```

```
rule execLoad(it matches tagged ELoad{dst:.rd,addr:.av});
    rf.upd(rd, dMem.read(av)); bu.deq();
endrule
```

```
rule execStore(it matches tagged EStore{value:.vv,addr:.av});
    dMem.write(av, vv); bu.deq();
endrule
```

- ◆ execLoad < fetch ?
 - Same as execAdd, i.e.,
 - rf: upd < sub
 - bu: {first, deq} < {find, enq}
- ◆ execStore < fetch ?
 - bu: {first, deq} < {find, enq}

March 5, 2008

<http://csg.csail.mit.edu/6.375>

L11-10

Properties Required of Register File & FIFO to meet performance specs

◆ Register File:

- $rf.upd < rf.sub$

◆ FIFO

- $bu: \{first, deq\} < \{find, enq\} \Rightarrow$
 - ◆ $bu.first < bu.find$
 - ◆ $bu.first < bu.enq$
 - ◆ $bu.deq < bu.find$
 - ◆ $bu.deq < bu.enq$

March 5, 2008

<http://csg.csail.mit.edu/6.375>

L11-11

The good news ...

- ◆ It is always possible to transform your design to meet desired concurrency and functionality
 - Though critical path and hence the clock period may increase

March 5, 2008


<http://csg.csail.mit.edu/6.375>

L11-12

Register Interfaces

$read < write$

$write < read ?$



$read'$ – returns the current state when *write is not enabled*
 $read'$ – returns the value being written if *write is enabled*

March 5, 2008


<http://csg.csail.mit.edu/6.375>

L11-13

Ephemeral History Register (EHR)

[Rosenband MEMOCODE'04]

$read^0 < write^0 < read^1 < write^1 < \dots$



$write^{i+1}$ takes precedence over $write^i$

March 5, 2008

<http://csg.csail.mit.edu/6.375>

L11-14

Transformation for Performance

execAdd < fetch
execBzTaken < fetch
execLoad < fetch
execStore < fetch

```
rule fetch_and_decode (!stallfunc1(instr, bu));  
    bu.enq1(newIt(instr, rf));  
    pc <= predIa;  
endrule  
  
rule execAdd  
    (it matches tagged EAdd{dst:.rd,src1:.va,src2:.vb});  
    rf.upd0(rd, va+vb); bu.deq0(); endrule  
rule execBzTaken(it matches tagged Bz {cond:.cv,addr:.av}  
    &&& (cv == 0));  
    pc <= av; bu.clear(); endrule  
rule execBzNotTaken(it matches tagged Bz {cond:.cv,addr:.av}  
    &&& !(cv == 0));  
    bu.deq0(); endrule  
rule execLoad(it matches tagged ELoad{dst:.rd,addr:.av});  
    rf.upd0(rd, dMem.read(av)); bu.deq0(); endrule  
rule execStore(it matches tagged EStore{value:.vv,addr:.av});  
    dMem.write(av, vv); bu.deq0(); endrule
```

March 5, 2008

<http://csg.csail.mit.edu/6.375>

L11-15

One Element FIFO *using EHRs*

```
module mkFIFO1 (FIFO#(t));  
    EHRReg2#(t) data <- mkEHRReg2U();  
    EHRReg2#(Bool) full <- mkEHRReg2(False);  
    method Action enq0(t x) if (!full.read0);  
        full.write0 <= True; data.write0 <= x;  
    endmethod  
    method Action deq0() if (full.read0);  
        full.write0 <= False;  
    endmethod  
    method t first0() if (full.read0);  
        return (data.read0);  
    endmethod  
    method Action clear0();  
        full.write0 <= False;  
    endmethod  
endmodule  
    method Action enq1(t x) if (!full.read1);  
        full.write1 <= True; data.write1 <= x;  
    endmethod
```

March 5, 2008

<http://csg.csail.mit.edu/6.375>

L11-16

Experiments in scheduling


Dan Rosenband, ICCAD 2005

- What happens if the user specifies:

$Wb < Wb < Mem < Mem < Exe < Exe < Dec < Dec < IF < IF$

No change in rules

a superscalar processor!



Executing 2 instructions per cycle requires more resources but is functionally equivalent to the original design

March 5, 2008

<http://csg.csail.mit.edu/6.375>

L11-17

4-Stage Processor Results

Design	Benchmark (cycles)	Area 10ns (μm^2)	Timing 10ns (ns)	Area 2ns (μm^2)	Timing 2ns (ns)
1 element fifo:					
No Spec	18525	24762	5.85	26632	2.00
Spec 1	11115	25094	6.83	33360	2.00
Spec 2	11115	25264	6.78	34099	2.04
2 element fifo:					
No Spec.	18525	32240	7.38	39033	2.00
Spec 1	11115	32535	8.38	47084	2.63
Spec 2	7410	45296	9.99	62649	4.72

benchmark: a program containing additions / jumps / load's

Dan Rosenband & Arvind 2004

March 5, 2008

<http://csg.csail.mit.edu/6.375>

L11-18

Summary

- ◆ For most designs BSV Compiler does good scheduling of rules with some user annotations for priority
- ◆ However, for complex designs sometimes concurrency control is quite difficult and requires a good understanding on the part of the designer of the concurrency issues
- ◆ Performance specification is a good, safe solution but is not implemented in the compiler yet.
 - user can do manual “renaming” and use EHRs to meet most performance goals
- ◆ RWires can solve any problems but exacerbate the correctness issue
- ◆ Synchronous pipelines (single rule) can avoid many problems but is not recommended for complex designs